Assignment I - Financial Econometrics

Ronan Kenny (2691815)
Tommy Link (2691913)
Guus Bouwens (2701442)
Jip de Boer (2710065)

May 5, 2023

Part A: Volatility in the German stock market

Exercise 1

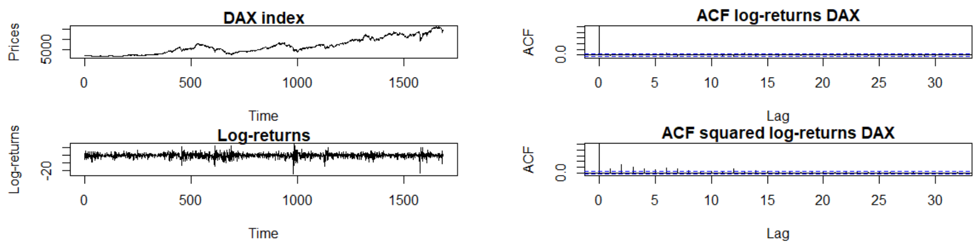

The plot of the DAX index and the corresponding log-returns will provide a visual representation of the evolution of the German stock market over time. The sample ACF for log-returns and squared log-returns may indicate whether there is autocorrelation in the data. It can also suggest the presence of volatility clustering or ARCH effects.
Based off the plots, the DAX index seems non-stationary while the log-returns seem to be stationary. We can make these conclusions based off the fact that the price level of the DAX index seems to change over time and the log-returns do the opposite, that is, they remain constant over time (around 0). The ACF plots will be discussed in the next section.

Exercise 2

The ACF plots for log-returns and squared log-returns will help assess the presence of autocorrelation or volatility clustering. Based on these plots, we can comment on the external consultant's statement regarding the relation between consecutive weekly returns. If there is no significant correlation, the consultant's statement may not be reliable. The log-returns show no significant autocorrelation across almost all lag frequencies (except 0, of course). There is however autocorrelation in the squared log-returns. This indicates that the external consultant's advice should not be followed since the past price changes do not reliably predict the future returns. More importantly, we can not necessarily show a positive return following a week with negative returns. In turn, we could actually counter the advice given by the consultant and say that the investment bank should not focus on the returns of the past, but rather the squared returns, since the squared returns will be able to give more information about the coming weeks due to the autocorrelation seen in the ACF above.

Exercise 3

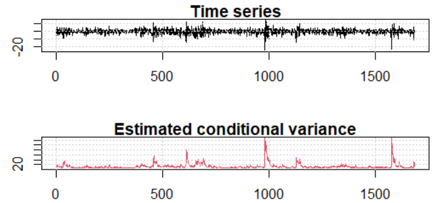

The GARCH(2,2) model parameter estimates provide insights into the volatility dynamics of the DAX log-returns. The filtered conditional variance plot above displays the evolution of volatility over time, highlighting periods of high or low market uncertainty. The parameter estimates for the GARCH(2,2) are $\hat{\theta} = (\hat{\omega}, \hat{\alpha_1}, \hat{\alpha_2}, \hat{\beta_1}, \hat{\beta_2}) = (0.865, 0.202, 0.108, 0.105, 0.512)$.

Exercise 4

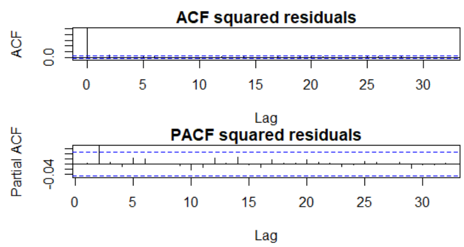

To address these colleagues' concerns about the GARCH(2,2) model specification, it would be helpful to analyze the model's residuals and conduct diagnostic tests, such as the Ljung-Box test, ARCH-LM test, or calculating the ACF and Partial ACF (PACF) of the squared residuals. As shown above, we look at the ACF and PACF of the squared residuals. We come to the conclusion that the plots show no autocorrelation in the residuals. This suggests that the model captures most of the relevant information in the data, and therefore provides a proper description of the data.

Exercise 5

Finding the best GARCH(p,q) model requires comparing different models by analyzing their goodness-of-fit, AIC/BIC values, and filtered variances. The "best" model should provide a good balance between model complexity and predictive power. We specifically look at the AIC / BIC values, as well as the paramters each model estimates by ML and the filtered variances. The great thing about looking at the AIC and BIC is that we can compare models with different amounts of parameters. Models with more parameters can sometimes "overfit" the model, but luckily the AIC and BIC can deal with overfitting by introducing a penalty term that specifically deals with the number of parameters in a model. We only compare GARCH(p,q) models for $p, q \leq 2$ because the AICs and BICs were increasingly worse for greater p's and q's. For the AIC and BIC, the best model has the lowest respective value. The best model based off the filtered variances (which all look very similar to one another) and having the lowest AIC and BIC is the GARCH(1,1) model.

GARCH(p,q)	AIC	BIC	$\hat{\theta}$
GARCH(1,1)	8186.6	8202.9	(0.54,0.19,0.75)
GARCH(1,2)	8189.1	8210.8	(0.55,0.21,0.10,0.64)
GARCH(2,1)	8188.8	8210.6	(0.59,0.18,0.03,0.75)
GARCH(2,2)	8190.2	8217.3	(0.84,0.20,0.10,0.47,0.16)

Table 1

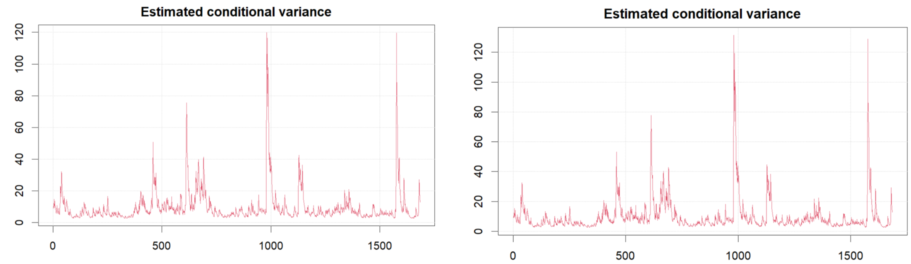

Filtered Volatility GARCH(1,1) and GARCH(1,2)

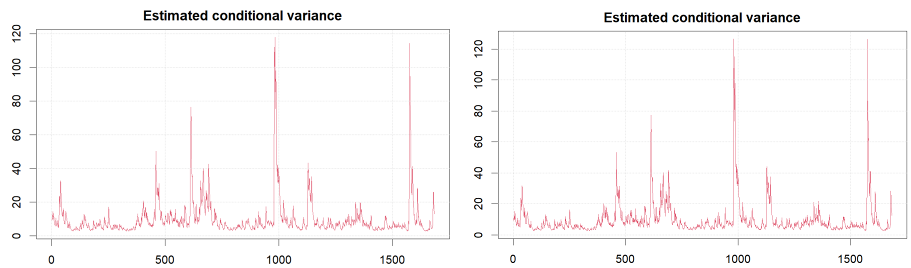

Filtered Volatility GARCH(2,1) and GARCH(2,2)

Exercise 6

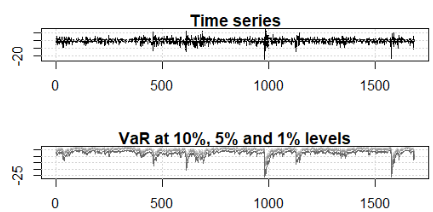

This statement is correct. Using the Jarque Bera Test (where the respective p-values are very close to zero and thus the null hypotheses of the errors and log-returns being normal are rejected), we can conclude that indeed the DAX log-returns and the errors do not follow a normal distribution. Obtaining the Value-at-Risk under the assumption of normal distribution of the error is inaccurate since the DAX log-returns are not normal. Assuming the normality of the errors will lead to overestimation of the risk because the normal distribution underestimates the probability of extreme events, and thus overeestimates the probability of moderate events. It is better to use a non-normal distribution, i.e., the student distribution, which would allow for extreme events more often.

Exercise 7

Estimating GJR-GARCH(1,1) model parameters and comparing them with the GARCH(1,1) model using AIC/BIC values helps determine which model is more suitable. If the GJR-GARCH(1,1) model captures the leverage effect and provides a better fit, it might be preferred over the GARCH(1,1) model. Our estimates for the GJR-GARCH(1,1) model are $\hat{\theta} = (0.739, 0.024, 0.748, 0.310)$ and the AIC and BIC are 8108.85 and 8130.57, respectively. Our estimates for the GARCH(1,1) model are $\hat{\theta} = (0.542, 0.191, 0.763)$ and the AIC and BIC are 8186.60 and 8202.89, respectively. Based of these values, we can determine the GJR-Garch(1,1) as the better model.

Part B: Portfolio management with multivariate GARCH models

Exercise 1

Our objective is to estimate a bivariate CCC model for the log-returns of DAX and HSI using the equation by equation approach. First, we estimate the univariate GARCH(1,1) for DAX log-returns and are left with the following estimates: $(\hat{\omega}, \hat{\alpha_1}, \hat{\beta_1}) = (0.542, 0.191, 0.763)$. Next, we estimate the univariate GARCH(1,1) for the HSI log-returns and get the following estimates: $(\hat{\omega}, \hat{\alpha_1}, \hat{\beta_1}) = (0.154, 0.077, 0.910)$.

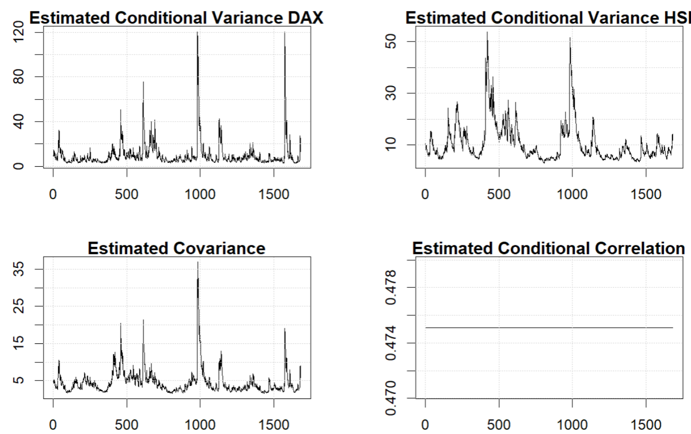

The CCC model is used to deal with the curse of dimensionality and cases where Σ_t might not be positive definite. At the expense of this, the conditional correlation matrix is constant (as seen in the plot of the estimated conditional correlation above). Despite this, we can see that the conditional covariance is still time varying.

Exercise 2

Below is the plot for the conditional variance and the α-VaR at a 1%, 5%, and 10% level for the portfolio of the bank (the darkest line is 1% and the lighest line is 10%).

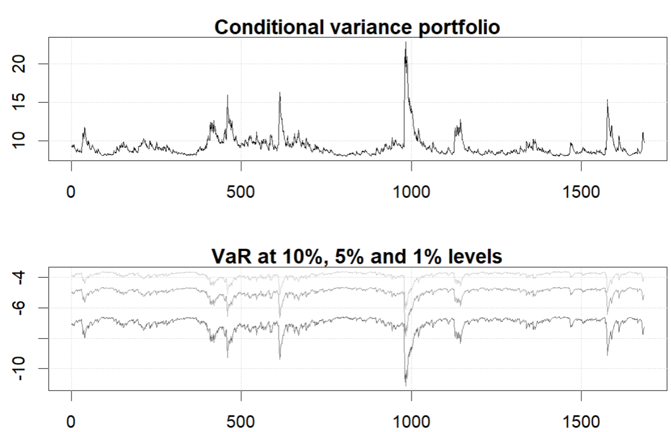

The conditional variance and α-VaR at 1% level for the bank's portfolio (70% in HSI and 30% in DAX) gives an idea of the risk associated with the current portfolio composition. The results show periods of high risk or low risk depending on the market conditions. We compare this to the different α levels (i.e., 5% and 10%) to compare the different percentage loss values the portfolio of the bank might have.

Exercise 3

The optimal weight at time $T+1$ is just the optimal weight at time T. These values for DAX and HSI are 0.472 and 0.528, respectively. The plot below shows the portfolio weights on DAX and HSI.

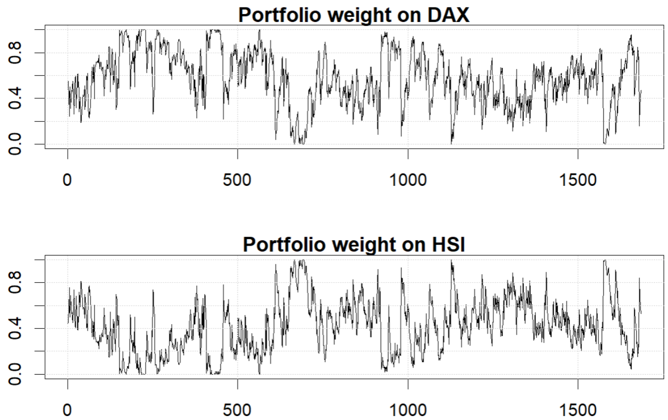

The optimal portfolio weights in terms of the Sharpe Ratio show how the weights should be adjusted over time to maximize the risk-adjusted returns. The optimal weight at time T+1 provides a suggestion for the bank on how to adjust the portfolio composition in the next period. The plot of the portfolio weights reveal how the weights have changed over time.

Exercise 4

The bivariate sDVECH model with covariance targeting will provide another perspective on the volatility dynamics between the two markets. Comparing the results with the CCC model, we observe differences in the estimated conditional variances, covariance, and correlation. The α-VaR at 1% level for the bank's portfolio (70% in HSI and 30% in DAX) slightly differ between the two models, indicating that the choice of the model can have an impact on the risk assessment. The parameter estimates for the model are $(\hat{\alpha}, \hat{\beta}) = (0.076, 0.895)$. The first picture below shows the plots for the estimated conditional variances, covariance, and correlation. The second picture shows the conditional variance portfolio and the VaR at the 10%, 5%, and 1% α levels.

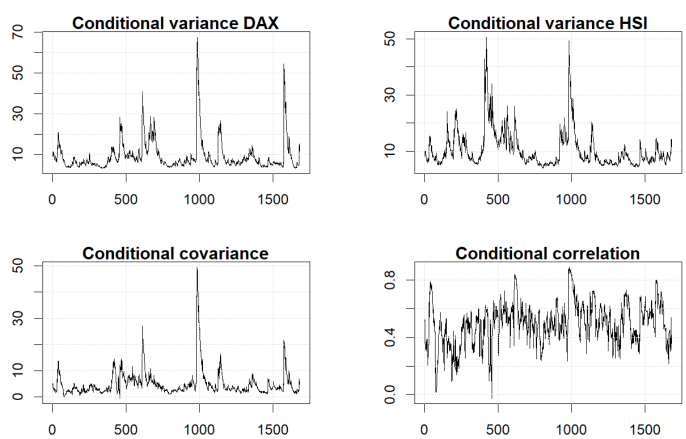

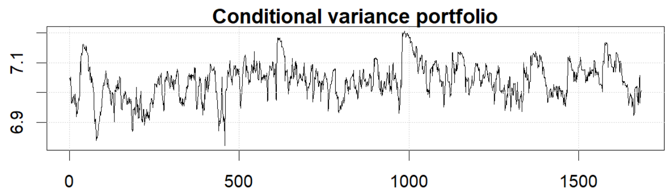

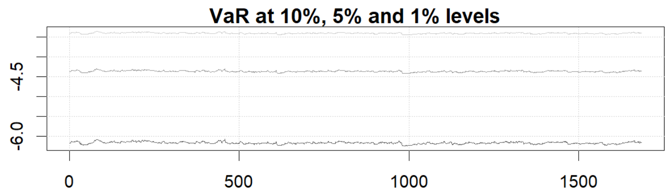

There is more obvious differences in the α-VaR percentage loss values above compared to Question 2 for 5% and 10% than 1%. The percentage loss values fluctuate far less for the sDVECH model than the CCC model. Perhaps this is down to the conditional correlation not being held constant in sDVECH model as it is in the CCC model.

Exercise 5

The forecasts of the volatility of the bank's portfolio for the next 52 weeks will give a sense of the expected risk associated with the portfolio in the future. The plot of the forecasts reveal periods of increased or decreased risk in the upcoming weeks. The variance forecast $\Sigma_T^2(h)$ in the sDVECH model is given by:

$$\Sigma_T^2(h) = Var(y_{T+h}^2|Y^T) = E(\Sigma_{T+h}^2|Y^T) \tag{1}$$

For $h > 1$

$$\Sigma_T^2(h) = W + (\alpha + \beta)\Sigma_T^2(h-1) \tag{2}$$

Then the volatility forecast of the portfolio forecast can be calculated as follows:

$$\sigma_{p,T}^2(h) = k^T \Sigma_T^2(h) k \tag{3}$$

k is a vector containing the constant portfolio weights, which in our case are 0.7 and 0.3. Below is the plot for the portfolio volatility forecast. We expect the portfolio volatility to decrease over the next 52 weeks.

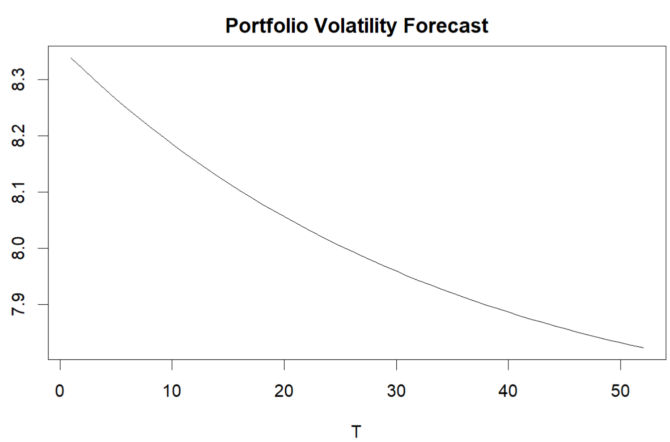

Exercise 6

By estimating a 3-dimensional CCC model for the three market returns (DAX, HSI, and S&P500), we can assess the relationship between all three markets. The estimated parameters and conditional covariance matrix plot provides insights into the correlations between the markets. The optimal portfolio at time T+1 for the three markets suggests how the bank should allocate its investments among the three markets to maximize risk-adjusted returns. The results show that investing in the US market (S&P500) could help further diversify the portfolio and reduce risk. The optimal portfolio weights at time T+1 are 0, 0.425, and 0.575 for DAX, HSI, and S&P500, respectively.

log-returns	$\hat{\theta} = (\hat{\omega}, \hat{\alpha}, \hat{\beta})$
DAX	(0.542,0.190,0.762)
HSI	(0.154,0.077,0.910)
S&P500	(0.549,0.192,0.761)

GARCH(1,1) Univariate Estimates

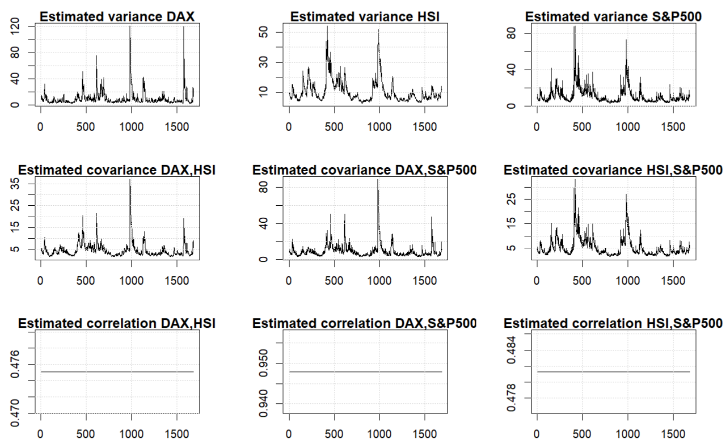

Plots of Estimated Variances, Covariance Matrices, and Correlations

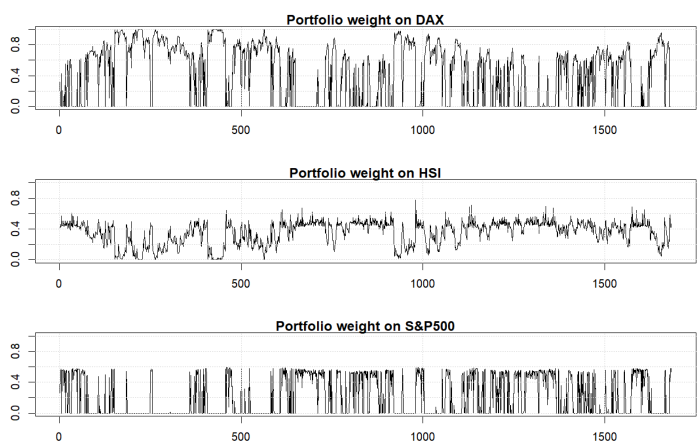

Portfolio Weights on the Three Market Returns